

013-32
53216
p-a

N92-14250

A 640-MHz 32-Megachannel Real-Time Polyphase-FFT Spectrum Analyzer

G. A. Zimmerman, M. F. Garyantes, M. J. Grimm
Communications Systems Research Section

B. Charny
Spacecraft Telecommunications Equipment Section

A polyphase-fast-Fourier-transform (FFT) spectrum analyzer being designed for NASA's Search for Extraterrestrial Intelligence (SETI) Sky Survey at the Jet Propulsion Laboratory is described. By replacing the time-domain multiplicative window preprocessing with polyphase filter processing, much of the processing loss of windowed FFTs can be eliminated. Polyphase coefficient memory costs are minimized by effective use of run-length compression. Finite word length effects are analyzed, producing a balanced system with 8-bit inputs, 16-bit fixed-point polyphase arithmetic, and 24-bit fixed-point FFT arithmetic. Fixed-point renormalization midway through the computation is seen to be naturally accommodated by the matrix FFT algorithm proposed. Simulation results validate the finite word length arithmetic analysis and the renormalization technique.

I. Introduction

A 2^{25} (33,554,432) channel, 640-MHz-wide polyphase-fast-Fourier-transform (FFT) spectrum analyzer is being designed at the Jet Propulsion Laboratory for the Search for Extraterrestrial Intelligence (SETI) Sky Survey. This spectrum analyzer will be used to separate two 320-MHz-wide polarizations into channels approximately 20 Hz wide for input to SETI signal detection algorithms. Construction of a prototype windowed-FFT spectrum analyzer [1] with 40 MHz of bandwidth and 2^{21} (2,097,152) channels has recently been completed. The new spectrum analyzer

design, similar to the prototype machine in many respects, is functionally divided into eight identical 80-MHz, 4-megachannel, real-input polyphase-FFT filter banks, each implemented as a pipelined special-purpose hardware signal processor. The spectrum analyzer functions consist of polyphase preprocessing, a 4-megapoint matrix-algorithm FFT and trigonometric recombination ("real-adjust") to compute the positive half of an 8-megapoint real FFT.

Other than the increases in bandwidth and number of channels, the main architectural difference from the

prototype spectrum analyzer system is that the new design is a polyphase-FFT spectrum analyzer rather than a windowed-FFT spectrum analyzer. The advantages of polyphase-FFT spectrum analysis, as well as a review of the supporting theory, are presented in [2] and are only touched on here. Similarly, features that the spectrum analyzer has in common with the prototype system can be found in [1] and are not described in detail here.

The remainder of this article will be divided into the following four main sections: a general description of the spectrum analyzer, a description of the polyphase-FFT filter bank implementation, a discussion of finite word length effects and the fixed-point arithmetic implementation, and results from system simulation.

II. General Description

The 320-MHz dual-polarization system is divided into four 80-MHz subbands per polarization. Each of the 80-MHz subbands, or "slices," is identical to the others, and each can be operated independently as an 80-MHz spectrum analyzer. A functional block diagram of the 320-MHz dual-polarization system is shown in Fig. 1.

Like the prototype system, each 80-MHz slice of the spectrum analyzer is a pipelined architecture, allowing all stages of the polyphase-FFT algorithm to execute concurrently. As in the prototype system, a stage "bypass" capability and stimulus and response buffers provide built-in testability. Each 80-MHz slice of the system will have an 8-bit analog/digital (A/D) converter as its input, and 24-bit fixed-point arithmetic will be used for the FFT portion. The fixed-point arithmetic will be implemented in an application-specific integrated circuit (ASIC), jointly developed with the Telecommunications and Data Acquisition (TDA) Advanced Systems very large scale integration (VLSI) program. This saves both memory and arithmetic relative to a floating-point implementation. In addition, because all of the FFT is completely performed in 24-bit fixed-point arithmetic, the transforms before and after the matrix transposition, or "corner turn," can use identical boards, saving a unique board design. In contrast, the prototype performs the column FFTs in 16-bit fixed-point arithmetic and the row FFTs in 32-bit floating-point arithmetic.

Architectural improvements in the FFT portion of the new spectrum analyzer include replacing two spectrum-length (4-complex-megapoint) double buffers with single buffers and the removal of a third spectrum-length dou-

ble buffer.¹ This improvement significantly reduces the amount of memory in the spectrum analyzer, reducing both size and cost.

The user-loadable window function in the prototype system has been replaced by polyphase filter preprocessing. The polyphase preprocessor is capable of operating as a window function, because a windowed discrete Fourier transform (DFT) is a degenerate case of a polyphase DFT filter bank. For more information on polyphase DFT filter banks, see [2].

As in the prototype system, the FFT is implemented using a matrix-style DFT pipe. By decomposing the transform into shorter row and column transforms, the long-delay memories and coefficient ("twiddle factor") storage required for the 4-megapoint FFT are concentrated in a single matrix transposition and complex multiplication stage, making the actual FFT arithmetic boards much simpler. The spectrum analyzer performs a 4-megapoint (2^{22}) DFT as 4096 point column DFTs followed by 1024 point row DFTs with multiplication of the matrix entries by complex rotation factors between the row and column transforms. The row and column DFTs are each performed as pipelined radix-4 FFTs.

While the prototype system was implemented in wire-wrap technology, the greater bandwidth and the highly repetitive nature of the new system require implementation using multilayer printed circuit boards. Because the new system takes advantage of improvements in technology, while it is functionally quite similar to the prototype, the detailed hardware designs are all new. A functional block diagram is shown in Fig. 2.

III. Polyphase FFT Structure

Because the spectrum analyzer channels are considered independently by the signal detection algorithms, the worst case processing loss [3], that is, the maximum attenuation of a continuous-wave (CW) signal in a channel's passband (in dB) plus the ratio of a channel's equivalent noise bandwidth to the channel spacing (in dB), is crucial to the system's sensitivity to weak signals. This is rarely less than 3 dB for windowed FFTs [3]. This is a result of the fact that for a windowed FFT, the time aperture over which the signal is considered (in seconds) is exactly equal to the reciprocal of the FFT channel spacing (in hertz).

¹ R. Brown, *Using Single Buffers and Data Reorganization to Implement a Multi-Megasample FFT* (internal document), Jet Propulsion Laboratory, Pasadena, California.

Consider an FFT with channel spacing N_t times finer than the desired spacing. This FFT operates on a time aperture N_t times longer than the FFT that would provide the desired channel spacing. Call the shorter FFT time aperture N_1 samples and the longer FFT $N_1 N_t$ samples. Now, apply a time-domain multiplicative window with the desired low-pass-filter transfer function to the long FFT, compute the FFT, and discard all but every N_t th FFT channel, i.e., retain only those bins whose center frequencies correspond to a center frequency in the shorter length N_1 transform. Now shift the $N_1 N_t$ sample input vector by N_1 samples and repeat the procedure. This process is equivalent to a polyphase-FFT filter bank, implementing a bank of identically shaped finite impulse response (FIR) band-pass filters, each centered on the bin center frequencies of the shorter, N_1 -long FFT.

In summary, the polyphase-FFT filter bank operates on a time aperture larger than its resolution, allowing bin shapes to encompass any transfer function that can be implemented as an $N_t N_1$ tap FIR filter. With as few as 8 taps per polyphase branch, it is possible to reduce the worst case processing loss to less than 1 dB. In fact, by including the polyphase-filter preprocessing step prior to computing the FFT, as much as 2.6 to 2.9 dB can be gained in worst-case processing loss over Hanning or Blackman windowed FFTs. In large FFT systems such as this one, the polyphase preprocessing has a small computational cost when compared to the FFT. However, substantial gains over windowed-FFT techniques are made even for systems with 2- or 4-tap polyphase branches, making polyphase preprocessing more than appropriate for smaller spectrum analyzer systems. For more discussion of polyphase DFT filter banks, see [2].

The main cost incurred in the polyphase preprocessing is in memory. The polyphase preprocessing requires that $N_t N_1$ input points must be stored at any one time, whereas in the windowed FFT system only N_1 points of storage were required. This storage requirement is mitigated somewhat by the fact that the points to be stored are of the input word length that is typically much shorter than the FFT arithmetic word length. In addition, the polyphase coefficients may be quantized to only slightly longer than the input word length without significant loss. In our implementation, the input word length is 8 bits, and 12 bits is sufficient for the polyphase coefficients. Due to the fact that our desired bin transfer functions are based on ideal band-pass filters, the prototype low-pass polyphase transfer function does not deviate significantly from an ideal low-pass filter. Correspondingly, the polyphase coefficients do not deviate significantly from the Fourier transform of an ideal low-pass filter, the Sinc function. As a re-

sult, the maximum rate of change of coefficients is limited, and a minimum run length can be found for quantized coefficients. For 12-bit coefficient quantization, the minimum run length is slightly more than 2048 coefficients, allowing for effective use of run-length compression to minimize the coefficient memory cost.

The minimum run length is greater than the row (second) FFT length, guaranteeing that there will be at most one transition per row. As a result, the rows may be separately run-length compressed, simply by identifying two coefficient values and the location of the transition. This is advantageous in a pipelined signal processor, where the order of data points into the board may be scrambled both within columns and by columns, because it allows run-length coding to be efficiently implemented as shown in Fig. 3.

Hardware simplifications result by setting the real and imaginary polyphase coefficients equal. Since the FFT is to be performed with the real-adjust algorithm, the real samples represent even-time indices while the imaginary samples represent the odd-time indices. Constraining real and imaginary polyphase coefficients to be identical means that the impulse response of the prototype low-pass FIR filter is constrained to change values only on even samples. Such a constraint will alter the transfer function of the resulting filter. One can use a matrix formulation of the DFT ($N_t N_1$ by 2 real-only points) to examine the transfer function of the resulting filter. Since the impulse response changes value only on even samples, the initial two-column DFTs ($N_t N_1$) are identical, and, in fact, both are the desired low-pass prototype transfer functions centered on the zero-frequency bin. The relevant portions of the entire transfer function are computed by taking the two-point DFTs of the "twiddle-in-the-middle" complex rotation factors $\exp(-j2\pi[(\text{row no.} \times \text{col. no.})N_1])$ in the passband and near-transition-band regions of the low-pass initial DFT results. The resulting transfer function is mainly low pass, with a replicant passband, attenuated by 135 dB, centered at the Nyquist frequency (π). For a system with a 41-dB signal-to-noise ratio (SNR) CW input, this results in a spur 28 dB below the noise floor in the output, requiring about 360,000 accumulations or about 5.2 hours to reach a 0-dB bin output SNR. It is also important to note that this spur is at the limit of the 24-bit representation, since the power of the smallest 24-bit number ($2^{23} + 0j$) is 141.5 dB down from the largest power value ($-1 - j$), and, as such, is at about the same level as output arithmetic and quantization noise. A simulation was performed with the real and imaginary coefficients constrained to be identical, and the results validated the above analysis.

IV. Finite Word Length Effects

With 8-bit input quantization, 16-bit fixed-point polyphase arithmetic and 24-bit fixed-point FFT arithmetic provide balanced system performance, with the input quantization (A/D) noise dominating. The spectrum analyzer has been designed with 8-bit input quantization based on considerations of A/D converter technology and the maximum expected RFI levels.

Noise due to finite word lengths can be divided between the A/D converter and the spectrum analyzer arithmetic. A balanced system would have roughly equal contributions from each source. The purpose of the SETI system is to detect small signals in the presence of interference and noise. Arithmetic and quantization noise will degrade the system performance by effectively increasing the input noise level, decreasing the SNR. The resulting noise power out will be:

$$\begin{aligned}\sigma_{out}^2 &= \sigma_{in}^2 + \sigma_{NQ}^2 + \sigma_{Arith}^2 \\ &= \sigma_{in}^2 \left(1 + \frac{\sigma_{NQ}^2}{\sigma_{in}^2} \left(1 + \frac{\sigma_{Arith}^2}{\sigma_{NQ}^2} \right) \right)\end{aligned}$$

where σ_{in}^2 and σ_{out}^2 are the input and output noise power, and σ_{NQ}^2 and σ_{Arith}^2 are the quantization and arithmetic noise, respectively. It is apparent that the sensitivity loss is

$$\frac{\sigma_{out}^2}{\sigma_{in}^2} = 1 + \frac{\sigma_{NQ}^2}{\sigma_{in}^2} \left(1 + \frac{\sigma_{Arith}^2}{\sigma_{NQ}^2} \right)$$

To put the input noise and the two digital noise factors into the same units, a translation from rms noise (volts) to quantization levels must be used; however, given a fixed input quantizer word length, it is apparent that the loss will be controlled by the ratio of the arithmetic to quantization noise.

Using the techniques in [4], the noise due to fixed-point multiplication round offs present in an output can be computed. This computation assumes that all round-off noise sources are white, mutually uncorrelated, zero mean, and uniformly distributed with a maximum value of 1/2 of one least significant bit. A radix-4 decimation-in-frequency FFT implementation was assumed. Each radix-4 FFT stage scales the data by 1/4 to ensure no overflows occur. Blind overflow-protection scaling of the FFT stages

would result in few significant bits to represent uncorrupted, noise-only or weak signal outputs, so the 24-bit words must be automatically renormalized at some point in the computation. This renormalization occurs at the corner-turn memory. Making the usual assumption of white, zero-mean, uniformly distributed input quantizer noise, noise due to the input quantizer was computed. The resulting ratio of the arithmetic noise to input quantization noise is given in Table 1. These results were used to calculate losses for the various input SNRs examined in the next section.

Note that renormalizing the data amplifies all sources of noise prior to the renormalization. When the renormalization shift is small, the majority of the arithmetic noise is due to the final stages of arithmetic, after the renormalization. Larger renormalization shifts reduce the relative effect of the fixed-point arithmetic. Note that the amount of renormalization shift depends on the maximum peak in the spectrum. In most cases, strong interference will be wideband relative to a spectrum analyzer channel (19 Hz), allowing a significant (>1) renormalization shift to be used.

In a pipelined processor implementation, renormalization requires a buffer capable of holding the entire vector. The matrix FFT algorithm was chosen for the FFT implementation based on experience with the SETI prototype [1]. The initial reason for this choice was because the algorithm concentrates pipeline delay memories and coefficient memory in a central matrix transposition and vector multiplication stage. In fixed-point FFTs, another advantage becomes apparent: the matrix transposition buffer provides a natural location for the vector renormalization. Renormalization can therefore be performed at the corner-turn memory, prior to the twiddle multiplication stage, without the cost of an additional buffer.

V. Simulation Results

The resulting system was simulated on a SUN-4 computer workstation with the appropriate word lengths, validating the analytical results and the renormalization approach. The simulated system had the following features:

- (1) A 12-tap-per-branch polyphase filter with 12-bit coefficients.
- (2) 16-bit fixed-point multiply-accumulate arithmetic in the polyphase filter computation.
- (3) 24-bit fixed-point arithmetic in the FFT, twiddle multiply, and real-adjust sections.

- (4) 32-bit IEEE floating-point power calculation.
- (5) A radix-4 4096 point by radix-4 1024 point matrix style 4-megapoint complex FFT with real-adjust and automatic renormalization between the row and column transforms.
- (6) 8, 10, or 16 bits of quantization at the input (A/D conversion).

The simulations were performed to test the performance in the presence of a strong signal (75 percent full scale on the input quantizer), measuring both output SNR and two-tone dynamic range. The simulations demonstrated 90-dB, close-in, two-tone dynamic range, detecting a weak signal 90 dB down from a strong signal 10 FFT channels away.

A. Finite Word Length SNR Degradation

The measured SNR out was compared with the SNR that would be output if the input quantization and arithmetic were infinitely precise. The ideal, infinite precision SNR is given by:

$$SNR_{ideal\ out} (dB) = SNR_{in} (dB) + Gain_{FFT} (dB) \\ + Gain_{PP} (dB) - EQNB_{PP} (dB)$$

where $Gain_{FFT}$ is the gain due to the number of channels in the FFT (2^{22} channels = 66.23 dB), $Gain_{PP}$ is the gain of the polyphase-DFT filter transfer function, dependent on the signal's frequency within the resolved FFT channel (-0.268 dB at 0.4 bins offset), and $EQNB_{PP}$ is the equivalent noise bandwidth of the polyphase-DFT filter transfer function relative to an ideal bandpass filter (0.346 dB in this 12-tap case). For our test cases, the strong signal frequency was held constant at 10 percent of the Nyquist frequency, resulting in an offset of 0.4 bins from the center frequency of the target bin.

Losses were predicted using the conventional round-off and quantization noise models, assuming white noise from each source. The results of the simulations are given in Table 2.

At lower SNRs there is excellent agreement with theory. For 8-bit input quantization, the A/D converter noise and spurs dominate the arithmetic contribution, as predicted; they can be observed from the lower SNR 8-bit measurements and a comparison of the 41-dB and 47-dB SNR 8-bit predictions and measurements with the 10- and 16-bit predictions and measurements. It may be noted

that at high SNRs, the predictions deviate from the measurements. While the measured losses for 8-bit inputs exceed the predictions, those for 10- and 16-bit inputs are less than the predicted loss. At high input SNRs, the assumption of white quantization and computational noise sources is violated. This is particularly true for the 8-bit input quantization noise with greater than 40 dB input SNR [5]. Since the 8-bit case behaves differently than either the 10- or 16-bit cases at these SNR levels, the increased loss for 8-bit input at 41- and 47-dB input SNR is attributed to high SNR input quantizer effects. It is important to notice that at input SNRs above 40-dB (output SNRs above 106 dB), quantization spurs from an 8-bit input quantizer become noticeable in the output, defining the upper limit of the spectrum analyzer's useful range in the presence of strong CWs.

Identifying the cause of decreased loss at high SNRs requires a closer examination of the sources of quantization noise. As the SNR increases, the loss due to the 24-bit fixed-point arithmetic grows. Unlike the 8-bit input quantization cases, the loss with 16-bit input quantization is almost entirely due to the arithmetic, and with 10-bit input, a large portion (78 percent) is due to the arithmetic. Due to overflow-preventing attenuation in the FFT arithmetic, which scales the data by 1/4 each stage, the lion's share of the FFT noise contribution comes from the final stages of FFT and real-adjust arithmetic. With high SNR CW inputs, beyond the 40-dB design range of the instrument, a significant fraction of the output noise power underflows the 24-bit fixed-point precision and is mapped to zero. This is especially apparent with 60-dB input SNR, where 12.5 percent (one-eighth) of the output values are complex zeros, and an additional 45 percent have either a zero real or imaginary component. Since arithmetic noise is computed from multiplication round offs, and multiplication by zero is an exact operation, the result is a decreased number of noise sources when high SNR tones are input. As a result, the arithmetic noise can be reduced by as much as 57.4 percent, or 3.71 dB in the 60-dB input case. This is sufficient to account for the discrepancies observed.

B. Two-Tone Dynamic Range

Measurements of the system's close-in two-tone dynamic range were performed. The two-tone dynamic range is defined as the maximum ratio of strong to weak signal levels at which a weak signal can be detected. Detection was performed without accumulating, providing a lower bound on the actual two-tone dynamic range of the instrument. Two-tone tests confirmed that a tone 90 dB down from the strong signal and 10 bins away is detectable, with-

out accumulation, as the maximum spectral peak outside of ± 4 bins from the strong signal.

The minimum detectable small signal level, allowing accumulation, and hence the two-tone dynamic range are determined by the maximum spectral level outside of the immediate spectral neighborhood of the strong signal. The maximum spur-to-noise ratio defines the minimum detectable SNR. The minimum detectable signal must be only slightly greater than the maximum spur. Given the very large number of noise samples in the spectrum, the ratio of the maximum to the average of the sample set would be tightly constrained (11.6 dB to 12.4 dB, 10 percent to 90 percent probability) if the output were white Gaussian noise only and contained no spurs. Some measured values and the probabilities that they are due to white Gaussian noise alone are given in Table 3.

Measurements made with greater than 8-bit inputs indicate that the maximum spur levels observed were due to the input quantization. These measurements confirm that a system with 8-bit inputs would have a two-tone dynamic range of $41 \text{ dB} + \text{Gain}_{FFT} + \min(\text{Gain}_{PP}) - EQNB_{PP} - 13.0 \text{ dB} = 93.6 \text{ dB}$. Examples of detection of a weak

signal 80 dB down from a strong signal 10 channels away are shown in Fig. 4.

VI. Conclusions

This article has described the latest results in a continuing effort to build more sensitive and broader bandwidth multichannel spectrum analyzers [1,6]. The thrusts of the current effort, described in this article, have been in the areas of signal processing, machine architecture, and technology utilization. The introduction of the polyphase-FFT provides signal processing gain superior to windowed FFTs, nearly independent of frequency location within a bin. Improvements in machine architecture have allowed the use of fixed-point arithmetic and have saved cost in both memory and arithmetic parts. Through the effective utilization of advances in technology, the bandwidth and number of channels of a single processor unit have each been doubled. The signal processor will be capable of operation at the limits of its 8-bit A/D converter, allowing up to 41-dB input SNR with less than 1-dB loss in sensitivity and exhibiting greater than 90-dB two-tone dynamic range.

References

- [1] M. P. Quirk, H. C. Wilck, M. F. Garyantes, and M. J. Grimm, "A Wideband, High-Resolution Spectrum Analyzer," *TDA Progress Report 42-93*, vol. January-March 1988, Jet Propulsion Laboratory, Pasadena, California, pp. 188-198, May 15, 1988.
- [2] G. A. Zimmerman and S. Gulkis, "Polyphase-Discrete Fourier Transform Spectrum Analysis for the Search for Extraterrestrial Intelligence Sky Survey," *TDA Progress Report 42-107*, vol. July-September 1991, Jet Propulsion Laboratory, Pasadena, California, pp. 141-154, November 15, 1991.
- [3] F. J. Harris, "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform," *Proceedings of the IEEE*, vol. 66, no. 1, 1978.
- [4] A. V. Oppenheim and R. W. Schaffer, *Digital Signal Processing*, Chapter 9, Englewood Cliffs, New Jersey: Prentice-Hall, 1975.
- [5] M. J. Flanagan, "The Behavior of Quantization Spectra as a Function of Signal-to-Noise Ratio," *TDA Progress Report 42-107*, vol. July-September 1991, Jet Propulsion Laboratory, Pasadena, California, pp. 155-168, November 15, 1991.
- [6] G. A. Morris and H. C. Wilck, "JPL 2²⁰ Channel 300 MHz Bandwidth Digital Spectrum Analyzer," presented at the IEEE International Conference on Acoustics, Speech, and Signal Processing, Tulsa, Oklahoma, 1978.

Table 1. Arithmetic noise/Input quantization noise

Renormalization shift	Input quantization	
	8 bits, dB	10 bits, dB
0	-0.64	11.40
1	-6.66	5.39
2	-12.66	-0.62

Table 2. Simulation results

Input quantization, bits	SNR _{in} , dB	Ideal SNR _{out} , dB	Predicted loss, dB	Measured loss, dB
8	30	95.61	0.05	0.02
8	35	100.61	0.15	0.18
8	38	103.61	0.29	0.41
8	41	106.61	0.56	0.84
8	47	112.61	1.91	2.53
10	47	112.61	0.52	0.22
16	41	106.61	0.11	-0.065
16	60	125.61	4.70	1.415

Table 3. Maximum noise spur to average > 3500 bins from strong signal

Input, bits	SNR _{in} , dB	Strong signal/weak signal	Max spur/noise average, dB	Probability (max./avg. > x dB) (noise alone), percent
8	30	Single tone	11.60	90
8	35	Single tone	11.71	80
8	38	Single tone	11.54	>90
8	41	Single tone	12.80	2
10	47	Single tone	11.49	>90
8	41	80 dB	12.87	2
8	41	90 dB	13.00	0.9
16	41	80 dB	11.61	90
16	41	90 dB	11.49	>90

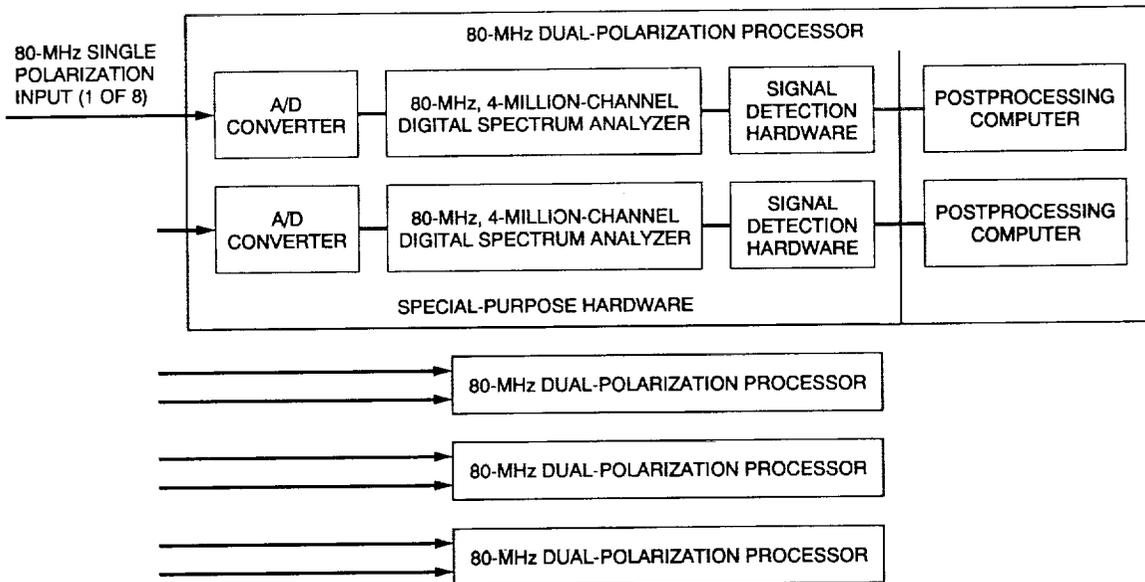


Fig. 1. SETI-MOP sky survey operational signal processor (320-MHz dual polarization).

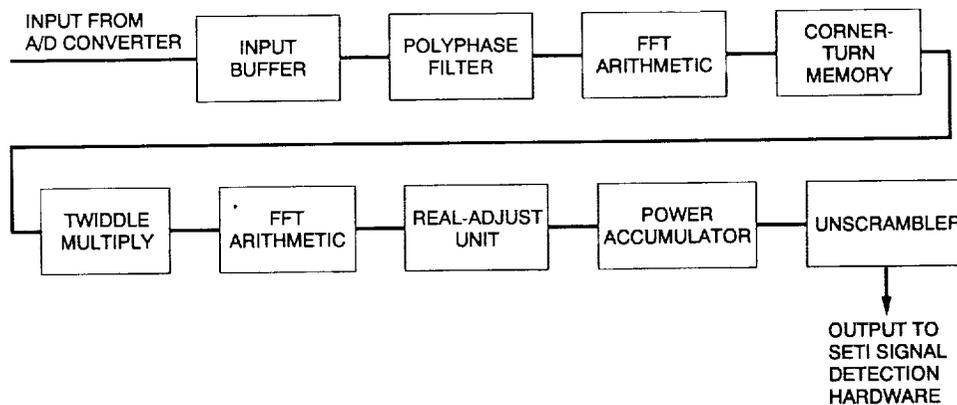


Fig. 2. An 80-MHz, 4-million-channel digital spectrum analyzer.

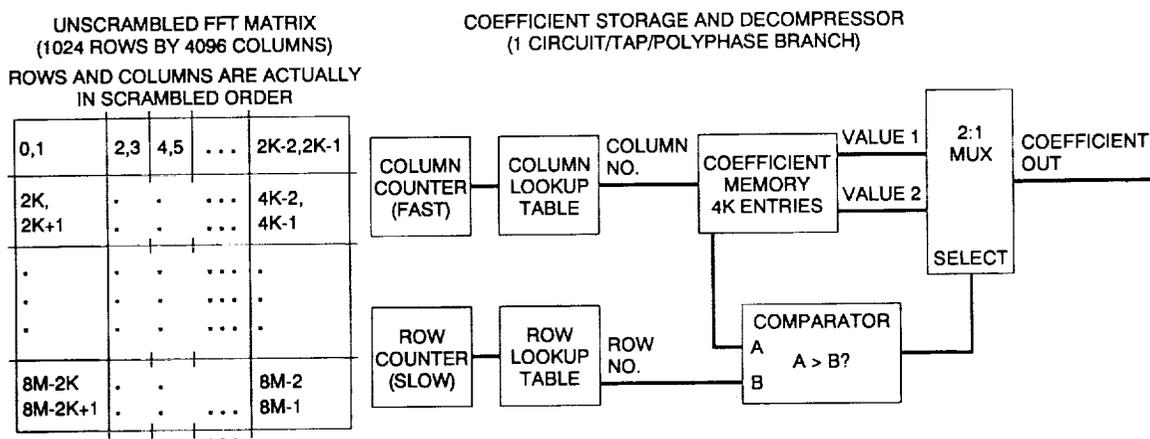


Fig. 3. Compression of polyphase filter coefficients.

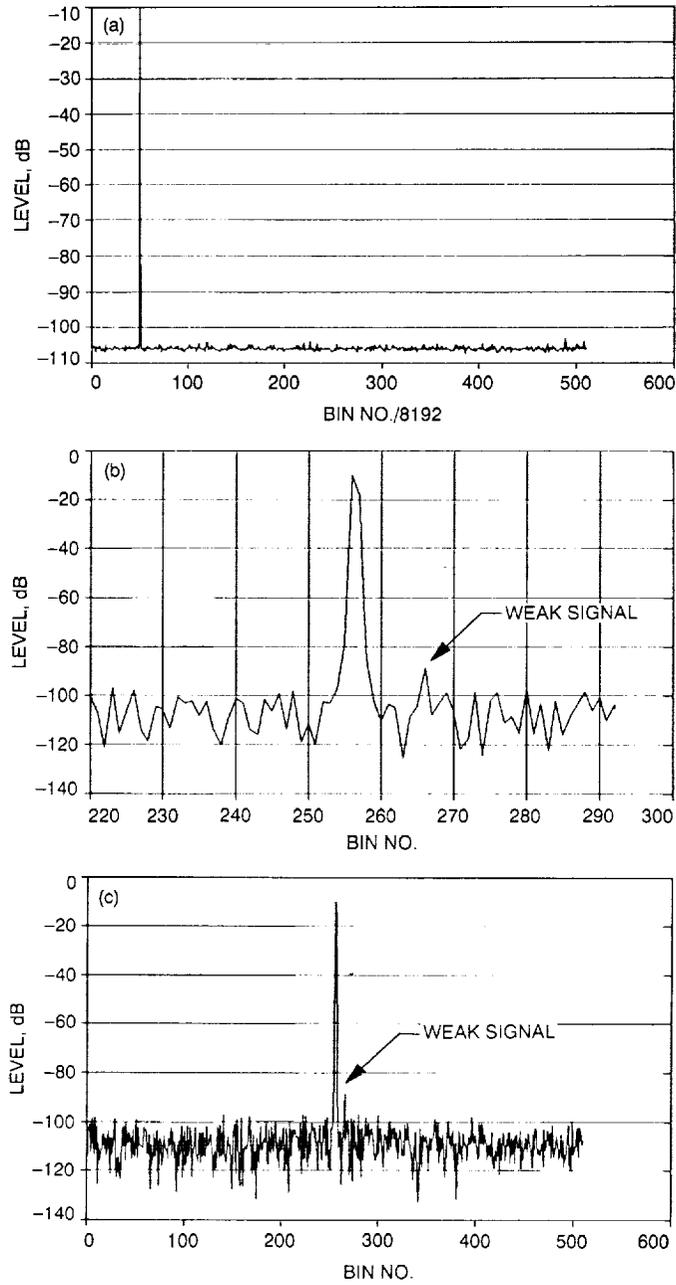


Fig. 4. Two-tone dynamic range test, 80-dB case: (a) full 4-megapoint spectrum with 8K compression by maximum; (b) signal region uncompressed; and (c) expanded signal region uncompressed.